

Map-making in small field modulated CMB polarisation experiments: approximating the maximum-likelihood method

D. Sutton¹, B.R. Johnson¹, M.L. Brown², P. Cabella^{1,3}, P.G. Ferreira¹, and K.M. Smith⁶

¹ *Oxford Astrophysics, University of Oxford, Denys Wilkinson Building, 1 Keble Road, Oxford OX1 3RH, United Kingdom*

² *Cavendish Laboratory, University of Cambridge, J J Thomson Avenue, Cambridge, CB3 0HE, United Kingdom*

³ *Universita di Roma 'Tor Vergata', Dipartimento di Fisica, via della ricerca scientifica 1, 00133 Rome, Italy*

⁴ *Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge. CB3 0HA, United Kingdom*

3 November 2008

ABSTRACT

Map-making presents a significant computational challenge to the next generation of kilopixel CMB polarisation experiments. Years worth of time ordered data (TOD) from thousands of detectors will need to be compressed into maps of the T, Q and U Stokes parameters. Fundamental to the science goal of these experiments, the observation of B-modes, is the ability to control noise and systematics. In this paper, we consider an alternative to the maximum-likelihood method, called *destriping*, where the noise is modelled as a set of discrete offset functions and then subtracted from the time-stream. We compare our destriping code (Descart: the DESTriping CARTographer) to a full maximum-likelihood map-maker, applying them to 200 Monte-Carlo simulations of time-ordered data from a ground based, partial-sky polarisation modulation experiment. In these simulations, the noise is dominated by either detector or atmospheric $1/f$ noise. Using prior information of the power spectrum of this noise, we produce destriped maps of T, Q and U which are negligibly different from optimal. The method does not filter the signal or bias the E or B-mode power spectra. Depending on the length of the destriping baseline, the method delivers between 5 and 22 times improvement in computation time over the maximum-likelihood algorithm. We find that, for the specific case of single detector maps, it is essential to destripe the atmospheric $1/f$ in order to detect B-modes, even though the Q and U signals are modulated by a half-wave plate spinning at 5-Hz.

Key words: cosmic microwave background – methods: data analysis – methods: statistical

1 INTRODUCTION

Measurements of the temperature and polarisation anisotropy of the cosmic microwave background (CMB) have been used to constrain a number of cosmological parameters to high precision (Smoot et al. 1992, Hanany et al. 2000, Masi et al. 2006, Hinshaw et al. 2008).

There are, however, still open questions that haven't been answered with the new results. In particular there is still no evidence for or against the existence of primordial gravity waves left over from a period of inflation at very early times. A bath of gravitational radiation should leave its imprint on the polarisation of the CMB in the form of a unique B-mode pattern. Such a signal will be faint compared to the polarisation that arises from density perturbations.

If we are to make a convincing measurement of B-modes (or rule out their existence), we need to construct an experiment with far greater sensitivity than the current state of the art. Given the physical limitations on the properties of individual detectors, the only realistic way of doing this is by observing the sky with multiple detectors for long periods of time. Most experiments that are being proposed have these characteristics.

A notable example is the C_ℓOVER experiment (North et al. 2008). This experiment will have 576 detectors and will focus on small fractions of the sky for periods of up to two years. With current planned technology, it is hoped that the noise per map pixel in the experiment will be below $0.8\text{-}\mu\text{K}$. To achieve this, it is essential to supplement the increased number of observations that multiple detector pixels bring by removing correlated noise in the map-making stage.

Solutions to the map-making problem aim to compress terabytes of time ordered data (TOD), acquired over months or years of scanning the sky, into a pixelised sky map, with noise uncorrelated between sky pixels. In each time ordered datum, the signal is typically swamped by noise, necessitating the process of binning and averaging at the heart of every map-making algorithm. Unfortunately, simple averaging is highly sub-optimal for this task because the noise in the TOD is correlated due to $1/f$ noise. One could attempt to high-pass filter the resulting low frequency $1/f$ noise drifts, thereby leaving random uncorrelated or *white* noise in the TOD. Such a method is not lossless, as it filters the signal as well. Whilst the effects of filtering can be mitigated in multipole space, by simulating its effects using Monte-Carlo simulations

(Hivon et al. 2002), the variance of the recovered power spectra are larger.

A number of previous experiments, including COBE, MAX-IMA and BOOMERANG, have opted to use the optimal maximum likelihood map-making algorithm (eg: Smoot et al. 1992, Hanany et al. 2000, Masi et al. 2006). This algorithm has been well discussed in the literature (eg: Tegmark (1997a), Natoli et al. (2001), Stompor et al. (2002), de Gasperis et al. (2005)) and is renowned for being slow and memory intensive. In addition, the algorithm requires accurate knowledge of the noise statistics of the TOD. Use of this algorithm is impossible for the next generation of high resolution, kilo-pixel, polarisation experiments, where we will require Monte-Carlo noise maps to reconstruct the noise covariance required for estimating B-modes on the cut sky (Smith & Zaldarriaga 2007).

Destriping is a promising alternative to the maximum likelihood approach (Keihänen et al. (2004), Keihänen et al. (2005)). By modelling the correlated noise in the TOD as a series of offset functions mapped onto the time-series, a nearly optimal map can be returned very quickly by solving for the offset function amplitudes. This system is much smaller than the full maximum likelihood system and can therefore be solved at a fraction of the computational cost. Destriping has been well discussed in the literature for the Planck experiment and favourably compared, in simulations, to the optimal solution: destriping can produce near-optimal maps in orders of magnitude less time (Poutanen et al. 2006, Ashdown et al. 2007a, Ashdown et al. 2007b). A variant of destriping has been successfully used in the Archeops balloon experiment analysis pipeline (Macias-Perez et al. 2007). However, these experiments are very different to upcoming ground-based B-mode experiments in terms of scanning strategy, presence of atmospheric noise and in the number of detectors.

Planck and Archeops use a *circular* scanning strategy: scanning the CMB in overlapping great circles before re-pointing to scan a new great circle. Scanning strategies are very different for ground based experiments, as a matter of necessity, because the scanning must be in azimuth only (constant elevation) in order to avoid extremely high atmospheric noise contamination from changing elevation. This causes a large difference in the level of cross-linking, which has a strong effect on the performance of map-making.

The biggest difference between ground based B-mode experiments and Planck or Archeops is the presence of atmosphere (Archeops flew so high it was essentially *above* the atmosphere). Atmospheric $1/f$ noise dominates the correlated noise in the TOD, even for the current generation of ground based experiments. It has a higher knee frequency and higher spectral index than the detector noise. For example, raw (undifferenced) TOD from the QUaD experiment (Hinderks et al. 2008) appears to exhibit $1/f^2$ atmospheric noise rather than simple $1/f$. Furthermore, the long term noise drifts will be very correlated between detector pixels (Bussmann, Holzapfel & Kuo 2005).

In this paper, we concentrate on experimental designs that modulate the faint polarisation signal to frequencies higher than the detector $1/f$ knee frequency f_{knee} , using a rotating half-wave plate, so that Q and U will be sampled in the white noise regime. Upcoming experiments with this design include C_lOVER (North et al. 2008), EBEX (Oxley et al. 2005) and SPIDER (MacTavish et al. 2007). These modulation experiments aim to mitigate possible detector $1/f$ using hardware. However, sophisticated map-makers will still be required to produce optimal Q and U maps. If we are to measure B-modes, we must treat the

systematics in the data properly. One of these systematics will be instrumental polarisation, which causes $T \rightarrow P$ leakage in the data (Johnson et al. 2007). If the leakage is significant, it will be vital to remove the leaked temperature signal, a process that will require a high-resolution, near-optimal map of the temperature anisotropies as observed by the instrument. Such leakage may also mix atmospheric $1/f$ noise into the Q and U data, for which a sophisticated map-maker like destriping will be needed to remove leaked $1/f$ in demodulated Q and U time-streams. The atmospheric noise itself may also contain a polarised component that would likewise have to be removed in the map-making (Hanany & Rosenkranz 2003).

This paper is the first in a series evaluating destriping as applied to ground-based, partial-sky, experiments. We examine the speed and accuracy of destriping for modulated experiments using simulations, dominated either by detector or atmospheric correlated noise, of a single detector pixel. This is akin to assuming that the noise is uncorrelated between detectors. We note that this is a poor assumption for the atmospheric $1/f$ component: we will generalise from this assumption in the next paper in the series (see also Appendix B).

In Section 2, we put our discussion on firm footing by deriving the optimal and destriping solutions to the map-making problems and discussing details of the implementation of the algorithms. In Section 3, we describe the simulation of the TOD and the scanning strategies. In Section 4, we first apply the algorithms to a preliminary signal-only simulation to test for bias and then proceed to apply the methods to the full Monte-Carlo simulations of TOD containing both signal and noise, comparing the performance of the algorithms directly with emphasis on speed and accuracy.

2 THE ALGORITHMS

Time ordered data, denoted by a time vector y_t , is formed from the sum of two time-space vectors, the sky signal S_t and total noise per observation n_t . The signal vector is created by the scanning of a telescope across the sky, which we consider to be innately pixelised. With this assumption, the signal stream becomes the result of the action of a $(N_{time} \times N_{pixel})$ projection operator, P_{tp} , on the sky map x_p .

This can be described in tensor notation by

$$y_t = P_{tp}x_p + n_t. \quad (1)$$

The reverse operation of P_{tp} is the binning operator $\mathbf{P}^T = P_{pt}$ (where the superscript T denotes the transpose), which sums the TOD into a map. The two operators acting together, $P_{pt}P_{tp}$, sum the *number of observations* per map pixel.

In the case of no noise, the map could be exactly returned by averaging: binning the TOD into a map and dividing each pixel by how many times it was observed

$$\vec{x}_p = (P_{pt'}P_{t'p})^{-1}P_{pt}\vec{y}_t, \quad (2)$$

which also returns the minimum-variance map in the case that the noise n_t is white. This defines the simplest “naive” map-making algorithm.

With polarisation measurements, we can measure 3 CMB skies: the unpolarised (or temperature) sky T and two skies corresponding to the Q and U polarisation parameters. In this case, the sky map vector becomes

$$\vec{x} = \begin{pmatrix} \vec{T} \\ \vec{Q} \\ \vec{U} \end{pmatrix}, \quad (3)$$

a vector of length $(3 \times N_{\text{pixel}})$, which has covariance matrix

$$\mathbf{C} = \langle x^T \vec{x} \rangle = \begin{pmatrix} \langle \mathbf{T}\mathbf{T} \rangle & \langle \mathbf{T}\mathbf{Q} \rangle & \langle \mathbf{T}\mathbf{U} \rangle \\ \langle \mathbf{Q}\mathbf{T} \rangle & \langle \mathbf{Q}\mathbf{Q} \rangle & \langle \mathbf{Q}\mathbf{U} \rangle \\ \langle \mathbf{U}\mathbf{T} \rangle & \langle \mathbf{U}\mathbf{Q} \rangle & \langle \mathbf{U}\mathbf{U} \rangle \end{pmatrix}. \quad (4)$$

with dimensions $(3N_{\text{pixel}} \times 3N_{\text{pixel}})$.

The projection operator \mathbf{P} describes the modulation of the Q and U signals into the TOD, expanding into T, Q and U polarisation pointing matrices

$$\mathbf{P} = \begin{pmatrix} \mathbf{P}_T \\ \mathbf{P}_Q \\ \mathbf{P}_U \end{pmatrix}. \quad (5)$$

The half-wave plate experimental design modulates the Q and U signals into the time stream at a frequency $\omega = 4f$, where the modulation frequency f is the frequency of half-wave plate rotation. The modulation frequency is chosen to be larger than the knee frequency of the correlated $1/f$ noise, such that the low intensity Q and U signals are sampled in the white noise regime, at a frequency $\omega > 4f_{\text{knee}}$.

The signal part of the modulated time stream is then

$$S_t = \frac{1}{2} P_{tp} (T_p + Q_p \cos 4\beta_t + U_p \sin 4\beta_t). \quad (6)$$

where β is the orientation angle of the half-wave plate with respect to the sky and P_{tp} is the simple pointing matrix from (1) (Johnson et al. 2007).

This gives us the polarisation pointing matrices from (5)

$$\begin{aligned} \mathbf{P}_T &= \frac{1}{2} P_{tp} \\ \mathbf{P}_Q &= P_{tp} \begin{pmatrix} \cos 4\beta \\ 2 \end{pmatrix} \\ \mathbf{P}_U &= P_{tp} \begin{pmatrix} \sin 4\beta \\ 2 \end{pmatrix}. \end{aligned} \quad (7)$$

2.1 The Maximum-Likelihood Solution to the Map-Making Problem

To estimate the sky map x_p , we begin with Bayes' theorem

$$P(\theta|DI) \propto P(D|\theta I) \times P(\theta|I), \quad (8)$$

where θ is our list of parameters, D is the data and I is all of the assumed background information of the experiment. In the map-making problem, the data is the TOD vector y_t and the parameters are the sky map x_p and the noise $C_{N,tt'} = \langle n(t)n(t') \rangle$. Bayes theorem now becomes

$$P(x_p C_{N,tt'} | y_t I) \propto P(y_t | x_p C_{N,tt'} I) \times P(x_t | I) P(C_{N,tt'} | I). \quad (9)$$

We assume that the noise is a stationary Gaussian realisation of $N(f)$, an underlying noise power spectrum, such that $\langle n_t \rangle = 0$, $C_{N,tt'} = C_N(t - t')$. Furthermore, we assume complete *a priori* knowledge of the noise power spectrum $N(f)$. This is like assuming a delta function prior on the noise, which effectively becomes part of the background information. With a flat prior on x_p , the posterior is simply proportional to the likelihood

$$P(\vec{x} | \vec{y} I) \propto \exp\left(-\frac{\chi^2}{2}\right), \quad (10)$$

with

$$\chi^2 = (\vec{y} - \mathbf{P}\vec{x})^T \mathbf{C}_N^{-1} (\vec{y} - \mathbf{P}\vec{x}) \quad (11)$$

where we have switched from tensor to matrix notation.

Solving for \vec{x} by minimising χ^2 , we find

$$\vec{x} = (\mathbf{P}^T \mathbf{C}_N^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{C}_N^{-1} \vec{y}, \quad (12)$$

which has pixel noise covariance matrix (obtained from the second derivative $\partial^2 \chi^2 / \partial \vec{x}^2$)

$$\mathbf{C} = (\mathbf{P}^T \mathbf{C}_N^{-1} \mathbf{P})^{-1}. \quad (13)$$

The matrix $(\mathbf{P}^T \mathbf{C}_N^{-1} \mathbf{P})$, sometimes called the *weight matrix*, is large and sparse. Traditional construction and inversion is prohibitive, with the memory requirement of storage scaling as $\mathcal{O}(N_{\text{pixel}}^2)$ and explicit inversion methods scaling as $\mathcal{O}(N_{\text{pixel}}^3)$. Fortunately, there are a number of computational tricks at our disposal.

The first trick is used to make the computation of \mathbf{C}_N^{-1} feasible. \mathbf{C}_N is an $(N_t \times N_t)$ matrix whose explicit construction and inversion is impossible. However, if we assume that the $1/f$ noise in the TOD is *stationary*, the ensemble average matrix \mathbf{C}_N can be well approximated as *circulant* (where each row is exactly the same as the row above, all shifted one column to the right). Systems involving circulant matrices can be inverted by deconvolution. If $\mathbf{A}x = b$, where \mathbf{A} is circulant, then $x = \mathcal{F}^{-1}[\mathcal{F}[b]/\mathcal{F}[A_{1n}]]$, where \mathcal{F} and \mathcal{F}^{-1} denote Fourier and inverse Fourier transformation respectively and A_{1n} is a vector formed of the first row of matrix \mathbf{A} .

This assumption is not valid for real data, as the TOD would have to wrap around to correlate the ends of the time-stream. For noise simulated by Fourier transforms however, the circulant approximation is exact. We note that for real data, the noise covariance is in fact a symmetric Toeplitz matrix, whose inverse can be calculated expensively in $\mathcal{O}(N_t^2)$ operations using Levinson's method (see eg: Press et al. 2002). The standard approach is to use the "MADCAP approximation" when inverting this matrix: the matrix is inverted as if it were circulant and then elements far from the diagonal are set to zero to enforce a Toeplitz structure, $C_N^{-1}(t - t') = 0$ for $|t - t'| > N_t/2$ (Borrill 1999). This approximate inverse covariance can then be convolved with the TOD using FFTs.

For the noise covariance of our simulations, we have that

$$\mathcal{F}[\mathbf{C}_{N1t}] = P(f), \quad (14)$$

where $P(f)$ is the ensemble average noise power spectrum and \mathbf{C}_{N1t} is the first row of \mathbf{C}_N .

The second trick speeds up the inversion of the weight matrix itself. Rather than inverting the matrix explicitly, we solve the system in (12) *iteratively*, using an algorithm from the conjugate gradients family. A number of maximum likelihood implementations have used a preconditioned conjugate gradients (pcg) algorithm for this task (eg: ROMA (Natoli et al. 2001), MapCUMBA (Dore et al. 2001) and SANEPIC (Patanchon et al. 2007)). However, we implement a more robust algorithm from the conjugate gradients family: MINRES, an algorithm capable of solving any invertible symmetric matrix, regardless of positive definiteness (Barret et al. 2006). By definition, the matrix is symmetric and should be invertible if there is a sufficient dispersion of polarisation measurement angles for each pixel. The weight matrix should, analytically speaking, be positive definite. We implement MINRES to protect against numerical departures from positive definiteness that can propagate from inaccuracies in noise estimation (Natoli et al. 2001).

The speed of iterative inversion depends on the condition number of the matrix (equivalent to the absolute ratio of the largest and smallest eigenvalues of the matrix). A poorly conditioned system is not well posed for iterative inversion, so we insert a preconditioning step into the algorithm. The system $\mathbf{A}x = b$ can be preconditioned using *preconditioner matrix* K as follows

$$K_1^{-1}AK_2^{-1}(K_2\vec{x}) = K_1^{-1}\vec{b} \quad (15)$$

where the left and right hand preconditioners K_1 and K_2 are defined by $K_1K_2 = K$. The preconditioner should be chosen such that $K^{-1} \approx \mathbf{A}^{-1}$, but it should also be quick to compute. In practice, choosing K is a trade off between improvements in convergence and added computation time. The size of the system prohibits the use of a sophisticated preconditioner, such as an incomplete Cholesky factorisation, so we use the simple but effective point Jacobi preconditioner. This is simply the diagonal of the weight matrix, which we approximate as

$$K_{pp'} = \sum_t \sum_{t'=t-\lambda_C}^{t+\lambda_C} (P_{pt}C_{N,tt'}^{-1}P_{t'p'})\delta_{pp'}, \quad (16)$$

for a correlation length λ_C ($\delta_{pp'}$ is the Kronecker delta).

Iterative inversion has a further advantage. The explicit calculation of the weight matrix is not required, we only need to reproduce the action of its matrix multiplication on the sky map \vec{x} . This is easily done with the following algorithm (in which the weight matrix is factorised into its component operations)

$$(\mathbf{P}^T \mathbf{C}_N^{-1} \mathbf{P})\vec{x} = P_{tp}^T \mathcal{F}^{-1} \left[\frac{\mathcal{F}[P_{tp}\vec{x}]}{\mathcal{F}[\mathbf{C}_{N1t}]} \right], \quad (17)$$

a process of projection, deconvolution (using fast Fourier transforms) and binning. This final trick has decreased the memory requirement for the matrix from $\mathcal{O}(N_{pixel}^2)$ to $\mathcal{O}(N_t)$.

It should be noted that the computation time for the optimal solution is still unfeasible for future data sets, despite the tricks implemented into the algorithm. The operational scaling of the algorithm is $\mathcal{O}(N_{it} \times N_t(1 + \log_2 |N_t|))$, where N_{it} is the number of iterations. The binning and projection operators (P_{tp}^T and P_{tp}) scale as $\mathcal{O}(N_t)$ so that each iteration is dominated by the FFT and inverse FFT (both of which scale as $\mathcal{O}(N_t \log_2 |N_t|)$).

2.2 An alternative: the destriping method

There is an alternative, approximate, method for modelling the noise in the time series. The noise vector \vec{n} can be modelled as uncorrelated white noise plus a series of discrete offsets that represent correlated noise. The amplitudes of the offsets are estimated and subtracted, as illustrated in Figure 1, whilst leaving the signal untouched. This approach, called destriping, has recently been derived in a maximum-likelihood context (Keihanen et al. 2004, Keihanen et al. 2005), whose notation we follow. The latter of these papers describes the MADAM algorithm, whose solution we use here.

We write the noise vector n_t as a sum of vectors of uncorrelated (white) noise, n_W , and correlated noise, n_{corr} ,

$$n_t = n_W + n_{corr}. \quad (18)$$

The essence of the destriping method is to approximate n_{corr} as an expansion of a set of offset functions, \mathbf{F}_α , with amplitudes a

$$n_{corr} = \sum_i \mathbf{F}_{ti} a_i. \quad (19)$$




Figure 1. *Upper panel* Plot of a section of raw TOD with long term $1/f$ noise drifts and the maximum likelihood offset functions (light grey) from Descart projected onto it. *Lower panel* Same as above except the TOD is signal only, so the offset amplitudes are all zero.

The simplest choice of offset function is a constant:

$$\mathbf{F}_{ti} = \begin{cases} 1 & t \in \Delta_i \\ 0 & \text{otherwise,} \end{cases} \quad (20)$$

where Δ_i is a chunk of the TOD. In this scheme, $F_i a_i$ is a discrete jump in TOD space from 0 outside chunk i to constant a_i inside chunk i . Thus, the sum in (19) produces a time domain vector of constant offsets with amplitudes \vec{a} approximating the correlated noise. More complicated offset functions can be chosen, for example Fourier series or Legendre polynomials.

The amplitudes \vec{a} will be Gaussian random numbers satisfying

$$\langle a_\alpha \rangle = 0 \quad (21)$$

$$\langle a_\alpha a_{\alpha'}^T \rangle = \mathbf{C}_{\mathbf{a}\alpha\alpha'}. \quad (22)$$

With this approximation for correlated noise, equation (1) for the TOD can be re-written as

$$\vec{y} = \mathbf{P}\vec{x} + \mathbf{F}\vec{a} + \vec{n}_W. \quad (23)$$

Again we minimise the χ^2 , but this time we require solutions for both the map \vec{x} and the amplitude vector \vec{a} . With \vec{a} as a second parameter, the likelihood of the TOD becomes

$$P(\vec{y}) = P(\vec{y}|\vec{x}, \vec{a})P(\vec{a}|\vec{x})P(\vec{x}). \quad (24)$$

However, the probability distribution of the amplitudes is in-

dependent of the CMB, which we consider to be deterministic and have no associated probability distribution. Thus, the amplitude probability distribution becomes an independent prior on \vec{a} , $P(\vec{a}|\vec{x}) = P(\vec{a})$, and the CMB prior is a constant. The likelihood is now given by

$$P(\vec{y}) = P(\vec{y}|\vec{x}, \vec{a})P(\vec{a}). \quad (25)$$

The likelihood in (25) is simply the white noise distribution, which is a Gaussian with covariance $\mathbf{N} = \sigma^2 \delta_{ij}$

$$P(\vec{y}|\vec{x}, \vec{a}) = (2\pi|\mathbf{N}|)^{-1/2} \exp\left(-\frac{1}{2}n_w^T \mathbf{N}^{-1} n_w\right). \quad (26)$$

The next step is to decide what prior information on the amplitudes to include through $P(\vec{a})$. If we have an estimate of the $1/f$ noise power spectrum, we can include prior information through the noise covariance matrix $\mathbf{C}_\mathbf{N}$. The probability distribution for the offsets is Gaussian

$$P(\vec{a}) = (2\pi|\mathbf{C}_\mathbf{a}|)^{-1/2} \exp\left(-\frac{1}{2}\vec{a}^T \mathbf{C}_\mathbf{a}^{-1} \vec{a}\right), \quad (27)$$

where $\mathbf{C}_\mathbf{a}$ is a re-projection of $\mathbf{C}_\mathbf{N}$ through $\mathbf{C}_\mathbf{a} = (\mathbf{F}^T \mathbf{F})^{-2} \mathbf{F}^T \mathbf{C}_\mathbf{N} \mathbf{F}$, which for our simulations is circulant (but for real data is symmetric Toeplitz and would be inverted and convolved as described in Section 2.1). Forming a χ^2 from these distributions we obtain a function to minimise for the amplitudes and the map:

$$\begin{aligned} \chi^2 &= -2 \ln |P(\vec{y})| \\ &= (\vec{y} - \mathbf{F}\vec{a} - \mathbf{P}\vec{x})^T \mathbf{N}^{-1} (\vec{y} - \mathbf{F}\vec{a} - \mathbf{P}\vec{x}) \\ &\quad + \vec{a}^T \mathbf{C}_\mathbf{a}^{-1} \vec{a}. \end{aligned} \quad (28)$$

The χ^2 can be simplified by writing the map \vec{x} in terms of the data through $\vec{x} = (\mathbf{P}^T \mathbf{N}^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{N}^{-1} (\vec{y} - \mathbf{F}\vec{a})$. This equation can be recognised as simple naive binning, as the noise in the destriped time-stream $\vec{y} - \mathbf{F}\vec{a}$ is white $\mathbf{N} = \sigma^2 \delta_{ij}$.

We can now gather terms involving \mathbf{P} and \mathbf{P}^T into a single operator \mathbf{Z}

$$\chi^2 = (\vec{y} - \mathbf{F}\vec{a})^T \mathbf{Z}^T \mathbf{N}^{-1} \mathbf{Z} (\vec{y} - \mathbf{F}\vec{a}) + \vec{a}^T \mathbf{C}_\mathbf{a}^{-1} \vec{a}, \quad (29)$$

where

$$\mathbf{Z} = \mathbf{I} - \mathbf{P}(\mathbf{P}^T \mathbf{N}^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{N}^{-1} \quad (30)$$

and \mathbf{I} is the identity matrix.

The operator \mathbf{Z} is a signal cleaning operator. The operation $\mathbf{Z}\vec{y}$ removes the signal component from \vec{y} by subtracting a naive map of \vec{y} projected onto TOD space. \mathbf{Z} also has the property $\mathbf{Z}^T \mathbf{N}^{-1} \mathbf{Z} = \mathbf{N}^{-1} \mathbf{Z}$.

We obtain an estimator for \vec{a} by minimising (29) with respect to \vec{a}

$$(\mathbf{F}^T \mathbf{N}^{-1} \mathbf{Z} \mathbf{F} + \mathbf{C}_\mathbf{a}^{-1}) \vec{a} = \mathbf{F}^T \mathbf{N}^{-1} \mathbf{Z} \vec{y}. \quad (31)$$

This is an inverse problem like that of the maximum likelihood algorithm. We can use many of the same tricks in solving it. We solve the system iteratively, using a preconditioned conjugate gradients (pcg) algorithm. The system is typically smaller than that in the previous section because $N_{chunks} < N_{pixel}$ for a single day. Again, we do not explicitly construct the matrix, rather we do each operation on both sides of (31) individually.

This system solves much more quickly than the maximum likelihood system. The operations on the left hand side of (31) scale

as $\mathcal{O}(N_t)$. The exception to this is the inversion of the circulant offset covariance matrix $\mathbf{C}_\mathbf{a}$, which can be achieved through cyclic deconvolution scaling as $\mathcal{O}(n_a \log_2 |n_a|)$. Each iteration of the destriping algorithm scales as $\mathcal{O}(N_t + n_a \log_2 |n_a|)$, in comparison to $\mathcal{O}(N_t(1 + \log_2 |N_t|))$, the iterative scaling of the maximum likelihood algorithm.

The system is preconditioned using preconditioner K such that $K^{-1} A \vec{x} = K^{-1} \vec{b}$, where

$$K = \mathbf{F}^T \mathbf{N}^{-1} \mathbf{F} + \mathbf{C}_\mathbf{a}^{-1} \quad (32)$$

is a circulant matrix and is inverted using cyclic deconvolution.

With the amplitude vector found, one subtracts the correlated noise approximation $\mathbf{F}\vec{a}$ from the TOD \vec{y} and naively bins the cleaned TOD to return the destriped map

$$\vec{x} = (\mathbf{P}^T \mathbf{P})^{-1} \mathbf{P}^T (\vec{y} - \mathbf{F}\vec{a}), \quad (33)$$

The noise covariance of this map will be

$$\mathbf{C} = (\mathbf{P}^T (\mathbf{N} + \mathbf{F} \mathbf{C}_\mathbf{a} \mathbf{F}^T)^{-1} \mathbf{P})^{-1}, \quad (34)$$

where $\mathbf{N} = \langle n_w n_w^T \rangle = \sigma^2 \delta_{ij}$ is the white noise covariance and $\mathbf{C}_\mathbf{a} = \langle a a^T \rangle$ is the covariance of the amplitudes.

3 SIMULATIONS

We produced 4 sets of simulations, each with 200 signal+noise and 200 further noise only realisations, from 12 hours of observing using a single detector sampling at $f_{sam} = 100$ -Hz for 2 scanning strategies. The simulations are summarised in Table 1. The experimental design was that of a modulation experiment scanning at $1^\circ s^{-1}$ with a half wave plate rotating at $f_{rot} = 5$ -Hz, corresponding to a polarisation modulation frequency $f_{mod} = 20$ -Hz.

Our aim is to examine the capabilities of different algorithms on the same data. For this goal, the white noise variance is essentially irrelevant: what we care about is the level of correlation between TOD and thence map pixels. To see B-modes will require years of observing with hundreds of detectors. For this work, in which we are simulating 12 hours of data from a single detector, we have chosen a white noise level that allows us to convincingly measure B-modes. For all our simulations, we use a heuristic $\text{NET} = 0.242 \mu K \sqrt{s}$.


Two noise scenarios are simulated. The first of these represents detector dominated $1/f$, as would be returned for a space-borne or balloon flight experiment (like MaxiPol (Johnson et al. 2007) or Archeops (Macias-Perez et al. 2007)) and has a $1/f$ power spectrum with spectral index $\alpha = 1.0$ and knee frequency 0.1-Hz, the projected noise correlation properties of the C_lOVER detectors.

The second scenario simulates $1/f$ dominated by atmospheric fluctuations, assumed to be un-polarised, which is expected to be the case for ground based experiments. The atmospheric $1/f$ noise is simulated with spectral index $\alpha = 1.9$ and knee frequency $f_{knee} = 0.2$ -Hz. We calculate these numbers from a rough fit to sample QUaD 100 GHz noise power spectra (see figure 36 of Hinderks et al. 2008) and they are fiducial: the level of atmospheric fluctuation depends upon wind and scanning velocities and evidence suggests it may vary between CMB observing sites (Bussmann, Holzapfel & Kuo 2005).

The TOD is simulated by taking a list of telescope pointings from a scanning strategy and pulling out the T, Q and U signals from simulated CMB maps at the right ascension and declination of

Table 1. Summary of the simulations

Strategy	tensor to scalar ratio	knee frequency	spectral index
Sabre	$r = 0.1$	$f_{knee} = 0.1\text{-Hz}$	$\alpha = 1.0$
Sabre	$r = 0.1$	$f_{knee} = 0.2\text{-Hz}$	$\alpha = 1.9$
Sabre	$r = 0.0$	$f_{knee} = 0.1\text{-Hz}$	$\alpha = 1.0$
Fence	$r = 0.1$	$f_{knee} = 0.1\text{-Hz}$	$\alpha = 1.0$

**Figure 2.** Input theoretical power spectra: the solid curve is C_l^{TT} ; the dashed curve is C_l^{EE} ; and the dot-dashed curve is C_l^{BB}

the pointing. The T, Q and U signals are combined into time streams using equation (6) and added to an instrumental noise stream.

3.1 CMB template

Two theoretical power spectra were calculated using the CAMB package¹ (Lewis, Challinor & Lasenby 2000) for a typical set of concordance parameters. The first spectrum was used to synthesise maps with B mode corresponding to an initial tensor to scalar ratio of $r = 0.1$ plus the expected B-modes from weak lensing. The T, E and B spectra from this model are shown in Figure 2. The second spectrum was calculated to produce zero B-mode simulations, using $r = 0$ and ignoring weak lensing B-modes.

The model power spectra were convolved with a fiducial symmetric Gaussian instrumental beam of FWHM 10 arc-minutes. From the initial spectra, 200 Gaussian realisations of the CMB were simulated using the HEALPix (Gorski et al. 2005) package *synfast*² using $n_{side} = 512$, giving resolution up to $l = 1024$.


3.2 Noise Simulation

The simulated noise stream is a sum of stationary Gaussian realisations of uncorrelated white noise and correlated $1/f$ noise, which together have power spectral density

$$P(f) = \frac{\sigma^2}{f_{sam}} \left(1 + \left(\frac{f_{knee}}{f} \right)^\alpha \right), \quad (35)$$

¹ <http://camb.info>

² <http://healpix.jpl.nasa.gov/html/facilitiesnode11.htm>

**Figure 3.** Ra/dec plot of the sabre scan pointing. The black curve is the rising field scan and the grey curve is the setting field scan.

where σ^2 is the white noise variance of a single observation, f_{sam} is the *sampling frequency* of the detectors (where integration time $t = 1/f_{sam}$), α is the spectral index and f_{knee} is the knee frequency of the spectrum.

The noise we have simulated is both stationary and Gaussian, satisfying

$$\langle n \rangle = 0 \quad (36)$$

$$\langle nn^T \rangle = \mathbf{C}_N. \quad (37)$$

Due to the stationarity property, the covariance of the noise between TOD becomes a function of time separation. The noise covariance matrix \mathbf{C}_N is circulant and symmetric, allowing us to quickly evaluate its inverse using a Fourier transform:

$$C_N^{-1}(t - t') = \left(\frac{2}{N_t} \right)^2 \int_{-\infty}^{\infty} P^{-1}(f) e^{-if(t-t')} df, \quad (38)$$

where $P^{-1}(f)$ is the inverse noise power spectrum.

3.3 Scanning strategies

We simulate TOD using two scanning strategies. The noise realisation in the TOD are the same for each scan, so this amounts to changing the signal in TOD. Scanning strategy has been shown to have an effect on the ability of map-makers to reduce noise correlation (Tegmark 1997b). Most important is the *connectivity* of the scan (or its degree of *cross-linking*). With greater cross-linking, map-makers should produce fewer stripes and smaller residuals. Whilst the isotropy of the scans is different, we have kept the same average integration time per pixel in each map.

Our scanning strategies are the following:

(i) *Sabre scan*: Simulated scanning strategy for ground based experiment sited in the Atacama desert. The scanning elevation is kept constant at 45° to minimise atmospheric noise. The telescope scans back and forth in azimuth with a sinusoidal velocity curve, whilst the field rises or sets through the scanning elevation. The field is observed twice per day as it rises and sets, resulting in a minimally cross-linked scan as shown in Figure 3. The experiment runs for ≈ 5 days to complete 12 hours of integration time.

(ii) *Fence scan*: For half the scanning time, the telescope slews back and forth horizontally with sinusoidal velocity whilst the field moves vertically through the scan at a much slower constant velocity. The scan is repeated for the second half of the scanning time

with the scanning direction and field drift directions swapped, producing a highly cross linked square field. In total, 12 hours integration is completed.

The sabre scan is representative of the typical level of cross-linking achievable from the ground when keeping the scanning elevation constant. This is a vital constraint, as if the scanning elevation is varied then noise sourced from changes in airmass swamps the faint CMB signal. Unfortunately, realistic scans are never ideally cross-linked. The fence scan represents a nearly ideally cross-linked scan that cannot be achieved from the ground, but it is included to probe the effects of cross-linking on the performance of destriping.

3.4 Diagnostics

We require from our diagnostics, comparable statistics indicating the fitness for purpose of the map-making algorithms. We use the following diagnostics:

(i) *RMS residual*: Root-mean-square of the residual map between the recovered maps and the input theoretical map used to generate the TOD.

(ii) *Residual angular power spectrum*: spatial information about the comparative temperature residuals is gained by analysing the angular power spectrum of the residual maps. The residual map is convolved with map field's window function W_l , so a modification of the MASTER method (Hivon et al. 2002) is used to return an unbiased binned estimate of the residual (noise) power spectrum.

(iii) *Pure pseudo- C_l E and B mode estimation*: The pure pseudo- C_l estimator (Smith 2006, Smith & Zaldarriaga 2007) is used to return estimates of the E and B mode angular power spectra C_l^{EE} and C_l^{BB} from the estimated Q and U maps. This estimator has been shown (Smith 2006) to return unbiased estimates of E and B without the E→B mixing from the ambiguous modes that arise from the field boundaries (Bunn 2002). The diagnostics are the mean estimates of the power spectra, the mean estimates of the E and B noise power spectra and the Monte-Carlo error bars for a single realisation. The details of the pure pseudo- C_l estimator are presented in Appendix A.

(iv) *Filter function*: The filter function F_l from a preliminary signal only simulation is used to analyse the signal error component to the estimated map residuals. The filter transfer function F_l (Hivon et al. 2002) gives the degree to which signal filtering by the map-making process affects the recovered power spectrum. If we don't want to lose any information on the signal then the imprint of filtering must be negligible ($F_l - 1 = 0$).

The power spectrum of the estimated map \tilde{C}_l is related to that of the input map C_l^I by (Natoli et al. (2001), Poutanen et al. (2006))

$$\langle \tilde{C}_l \rangle = F_l \langle C_l^I \rangle + \langle N_l \rangle \quad (39)$$

where N_l is the power spectrum of the noise bias. If the estimated map is made from signal only TOD, $\langle N_l \rangle = 0$ so F_l can be calculated by inverting (39).

Both \tilde{C}_l and C_l^I are convolved with the same beam and have the same pixelisation and sky mask. They have identical transfer matrices $K_{ll'}$ (see Appendix A), so we need not correct for mode coupling.

4 RESULTS

The simulations were analysed by three map-making algorithms: standard naive map-making (equation 2), optimal maximum likelihood map-making, and destriping. Destriping was repeated with various offset function baseline lengths (hereafter denoted by λ_C in units of time), ranging from a baseline length $\lambda_C = 1$ second through to $\lambda_C = 1000$ seconds (≈ 16 minutes). The destriping code, Descart, operated in two modes: *traditional* destriping mode (as in all papers prior to Keihänen et al. (2005)), and in a *covariant* destriping mode, making use of noise information through the offset covariance matrix C_a (Keihänen et al. 2005). The two modes are compared in Section 4.6. Elsewhere, all destriping is the covariant form.

Both the covariant destriping and maximum likelihood map-makers require prior noise information (through C_a and C_N^{-1} respectively). In reality, this noise information would have to be estimated from the TOD directly, highlighting the importance of the noise estimation step immediately prior to map-making. In this paper, we are not addressing the issues of noise estimation and so use ensemble average noise information through the power spectrum $P(f)$ used to generate the TOD noise. A number of approaches to noise estimation have been discussed in the literature (eg: Ferreira & Jaffe 2000, Amblard & Hamilton 2004).

4.1 Signal only maps

In addition to signal+noise simulations, an initial pure signal simulation was analysed first for the sabre strategy, using the same prior noise information as for the signal+noise simulations. In this limit, any departures in the estimated map from the input map are distortions of the signal caused by the map-making process itself. The filter functions for this simulation were calculated using (39) where C_l^I and \tilde{C}_l are raw pseudo- C_l s calculated by the HEALPix package *anafast*³ (Gorski et al. 2005).

Figure 4 shows the T, E and B filter functions for the maximum-likelihood (upper panel) and destriping (middle panel) algorithms. Destriping, covariant or otherwise, does not filter the signal at all, despite use of a noise prior through C_a (Figure 4, middle panel). The data term in the likelihood (29) forces the offset amplitude vector to the null vector, as the signal cleaning operator Z perfectly removes the signal from the fit.

The maximum likelihood algorithm does display some minimal signal distortion (Figure 4, upper panel). The magnitude of the distortion is effectively negligible, amounting to less than 0.001% of the signal at the worst multipoles ($l < 50$ for the T map). The filtering satisfies $\sigma_{\text{filtering}} \ll \sigma_{\text{noise}}$ at all multipoles and can be ignored.

This filtering can be explained through the presence of degenerate pixels. A minimum of three observations at different modulation angles β are required to reconstruct T, Q and U for each pixel. Any pixel for which this condition is not met is degenerate and must be ignored, else the problem becomes singular.

For Descart, the reconstruction inversion is conducted for each pixel separately - if any pixel's 3×3 pointing matrix is singular, then the pixel is irrecoverable. For the maximum-likelihood algorithm, the reconstruction is accomplished by the matrix inverted by the MINRES conjugate gradient inverter. If degenerate pixels are included, the MINRES iterations do not converge: the improvement

³ <http://healpix.jpl.nasa.gov/html/facilitiesnode5.htm>




Figure 4. Upper panel Filter function $F_l - 1$, for maximum likelihood algorithm showing minimal signal degradation from the initial *signal only* simulation. Middle panel $F_l - 1$ for destriping, displaying zero signal degradation. Lower panel Example filter function for naive map of TOD filtered with $C_N^{-1/2}$ filter, where the signals are distorted and T is heavily filtered

Table 2. Comparison of the performance of iterative inversion in the maximum-likelihood algorithm for different pixel exclusion conditions in the signal only simulation. N_{min} is the minimum number of hits required to accept a pixel whilst convergence is the rms residual upon exit of the MINRES algorithm. The $N_{min} = 2$ case includes degenerate pixels, so the matrix is singular and the iterations fail before the convergence criterion is reached.

N_{min}	#iterations	N_{pixel}	convergence
13	90	17499	10^{-6}
5	90	17520	10^{-6}
4	100	17524	10^{-6}
3	101	17531	10^{-6}
2	87*	17536	7.5^{-4}

* iterations failed

in residuals per iteration tends to zero before a reasonable convergence criterion is reached (such as 10^{-6} , Patanchon et al. 2007). If near-degenerate pixels remain, the condition number of the matrix increases, requiring more iterations to invert the matrix (see Table 2).

Such pixels will be common for ground based polarisation experiments, where the TOD is naturally split into day chunks and maps are made daily. It is well established that these pixels must be removed in such a way as to maintain time stream continuity (Stompor et al. 2002). Time ordered data sourced from a degenerate pixel is referenced to a junk pixel with zero signal outside of the estimated map that is ignored in the minimisation. Degenerate TOD are kept for the noise deconvolution step $C_N^{-1} \vec{y}$, where the TOD value is replaced by a constrained realisation of noise.

However, some errors from this method propagate through the deconvolution step of (17), which is applied approximately through FFTs, into neighbouring pixels along the scan direction, causing signal filtering.

This effect does not appear in the Descart destriped maps, as the degenerate TOD are removed from the binning/projection process of \mathbf{F}^T and \mathbf{F} respectively. There are no resultant gaps in the offset amplitude vector \vec{a} and so the evaluation of $C_a^{-1} \vec{a}$ using FFTs suffers no degradation due to discontinuities.

For comparison, the lower panel Figure 4 shows the filter function for high-pass filtered TOD using the $C_N^{-1/2}$ filter. In Fourier space, this filter is the square root of the noise filter in (38). It has the property that filtered TOD with $1/f$ noise have the same diagonal noise covariance matrix as TOD with white noise only, pre-whitening the TOD. The effect of the filter on the T signal is devastating. The effect on Q and U is mitigated here due to the modulation. Whilst the mean filter functions can be evaluated from Monte-Carlo simulations and then deconvolved from the power spectrum (Hivon et al. 2002), the wiggles in the functions are signal realisation dependent and will add to the variance of the recovered power spectrum. We suspect the high- ℓ bias in the B-mode filter function is due E→B mixing caused by distortion of the Q and U signals by the filtering. This will be studied in the next paper in the series.

The filtering effects here are separate from the signal filtering reported in Poutanen et al. (2006), Ashdown et al. (2007a) and Ashdown et al. (2007b). Our input CMB maps are at the same resolution (HEALPix $n_{side} = 512$) as the estimated maps, so the filtering error from applying FFTs to TOD including sub-pixel signal gradients is absent. We have considered the sky to be innately pixelised at the experimental resolution, which is one of the underlying




Figure 6. Variation of the mean dimensionless $\Delta_T = (\sigma/\sigma_{MLE})_T$ for 200 detector $f_{knee} = 0.1$ -Hz simulations with destriping function length λ_C in seconds. The diamond at $\lambda_C = 0.01$ -s is the maximum likelihood solution and the one at $\lambda_C = 43200$ -s is the naive solution of equation (2).

assumptions in map-making formalism. If the sky were pixelised as assumed, destriping would be lossless to machine precision.

Including sub-pixel gradients, it is expected that the maximum likelihood algorithm would filter the signal more than reported here: this is an avenue of on-going research.

4.2 Reconstructed Map Residuals

The residual map ϵ between a reconstructed map and the input simulated map can be defined for each of the Stokes parameters as

$$\begin{aligned} \epsilon_p &= x_p^{out} - x_p^{in} \\ &= \epsilon_p^S + \epsilon_p^N, \end{aligned} \quad (40)$$

where ϵ_p^S and ϵ_p^N represent signal error and pixel noise respectively, x_p^{out} is the recovered map estimate and x_p^{in} is the input map used to simulate the TOD. The signal error for the algorithms has been shown to be negligible, so $\epsilon_p \rightarrow \epsilon_p^N$.

4.2.1 Signal and detector noise simulations

We first look at the detector noise simulations with $f_{knee} = 0.1$. Differences between the maps from different algorithms are not perceptible by eye, as the map is dominated by the signal. However, differences are perceptible in the residual maps. The top row of Figure 5 contains residual T and Q maps for Descart (with $\lambda_C = 1$ -s), whilst the middle row show residuals from the same data mapped by the naive algorithm (in which the TOD is mapped using equation (2) with no pre-filtering). The characteristic $1/f$ noise stripes are visible in the naive T residuals and are notably absent in the Descart T residuals. The reduction in the correlated noise is best illustrated by the difference map between the destriped and naive maps (bottom left panel of Figure 5). The stripes in this map are the correlated $1/f$ noise in the pixel map that have been removed by the destriping process, visibly following the scanning pattern.

The Q and U residuals are very close white noise only due to the modulation of the signals out of the low f_{knee} detector $1/f$ so the destriped and naively mapped residuals (top and middle right

panels) are very similar. Their difference maps show a small change in the magnitude of the noise but no striping structure as for T. The structure in this difference map can be understood by noting that the $1/f$ noise that seeps through to Q and U maps is itself modulated by sinusoidal functions during the de-modulation process.

The maximum likelihood solution produces maps with the smallest possible residuals. The relevant statistic to measure the ability of the destriper to return maps cleaned of correlated $1/f$ noise is the ratio of the destriped map's root-mean-squared (rms) residual to that of the maximum likelihood map's rms residual, Δ_i :

$$\Delta_i = \left(\frac{\sigma}{\sigma_{MLE}} \right)_i, \quad (41)$$

where i is one of the Stokes parameters T, Q or U. This dimensionless statistic is independent of the white noise level in the map, it is the constant by which the noise in the map is multiplied above optimality.

Figure 6 shows the variation in mean (over 200 realisations) Δ_T with chunk length λ_C in the $f_{knee} = 0.1$ -Hz simulations. The unconnected diamonds are mean Δ_T for the maximum likelihood (corresponding to $\lambda_C = 0.01$ -s) and naive (corresponding to $\lambda_C = t_{int}$, the total integration time of the experiment) maps. As λ_C is decreased, Δ_T decays towards optimality at $\Delta_T = 1$. The best destriped maps are returned at the smallest chunks size considered ($\lambda_C = 1$ -s) and have only 0.4% higher pixel residuals than the optimal map, achieving 96% of the reduction in pixel residuals that the maximum likelihood method brings over naive binning. These numbers are relative to optimality and are dependent only on the correlated noise. They are independent of the magnitude of the white noise floor.

The destriping method models the noise as a white component + a correlated component described by a series of offset functions. The noise covariance C_N is modelled as

$$C_N \approx \sigma_W^2 \delta_{tt'} + FC_a F^T \quad (42)$$

where C_a is solely responsible for the off-diagonal part of C_N . When the resolution of C_a is increased, the approximation to the real C_N is more accurate and so the residuals are smaller. The residual reduction flattens as λ_C is decreased significantly below the noise correlation length (10 seconds for $f_{knee} = 0.1$ -Hz).

Mean Δ_Q for these simulations is shown in Figure 7. Despite the modulation of the Q and U signals to $f_{mod} = 20$ -Hz (200 times higher than the noise f_{knee}), there remains a gradient in the relative noise residuals between the optimal map and the naive map. Destriping converges to near optimal noise levels at $\lambda_C = 1$ -s, as for the T maps. The difference in pixel error between the algorithms is at most of order 0.1%.

The extra noise power in the naive T maps is at all angular scales, as shown in the residual angular power spectra in the upper panel of Figure 8. The curves in this plot are the mean T residual angular power spectra for all 200 realisations. The residual power spectrum for the optimal maximum-likelihood algorithm is the solid curve. The long dashed curve is from the Descart 1-s residuals, the dotted curve from the Descart 400-s residuals and the dot-dashed curve is from the naive residuals. The 400-s baseline Descart spectrum is included as an example of “quick and dirty” destriping as opposed to using destriping to replace the maximum-likelihood approach. This coarser noise model makes some improvement to noise at large scales but fails to reduce noise power at scales beneath its C_a resolution. The best 1-s baseline Descart

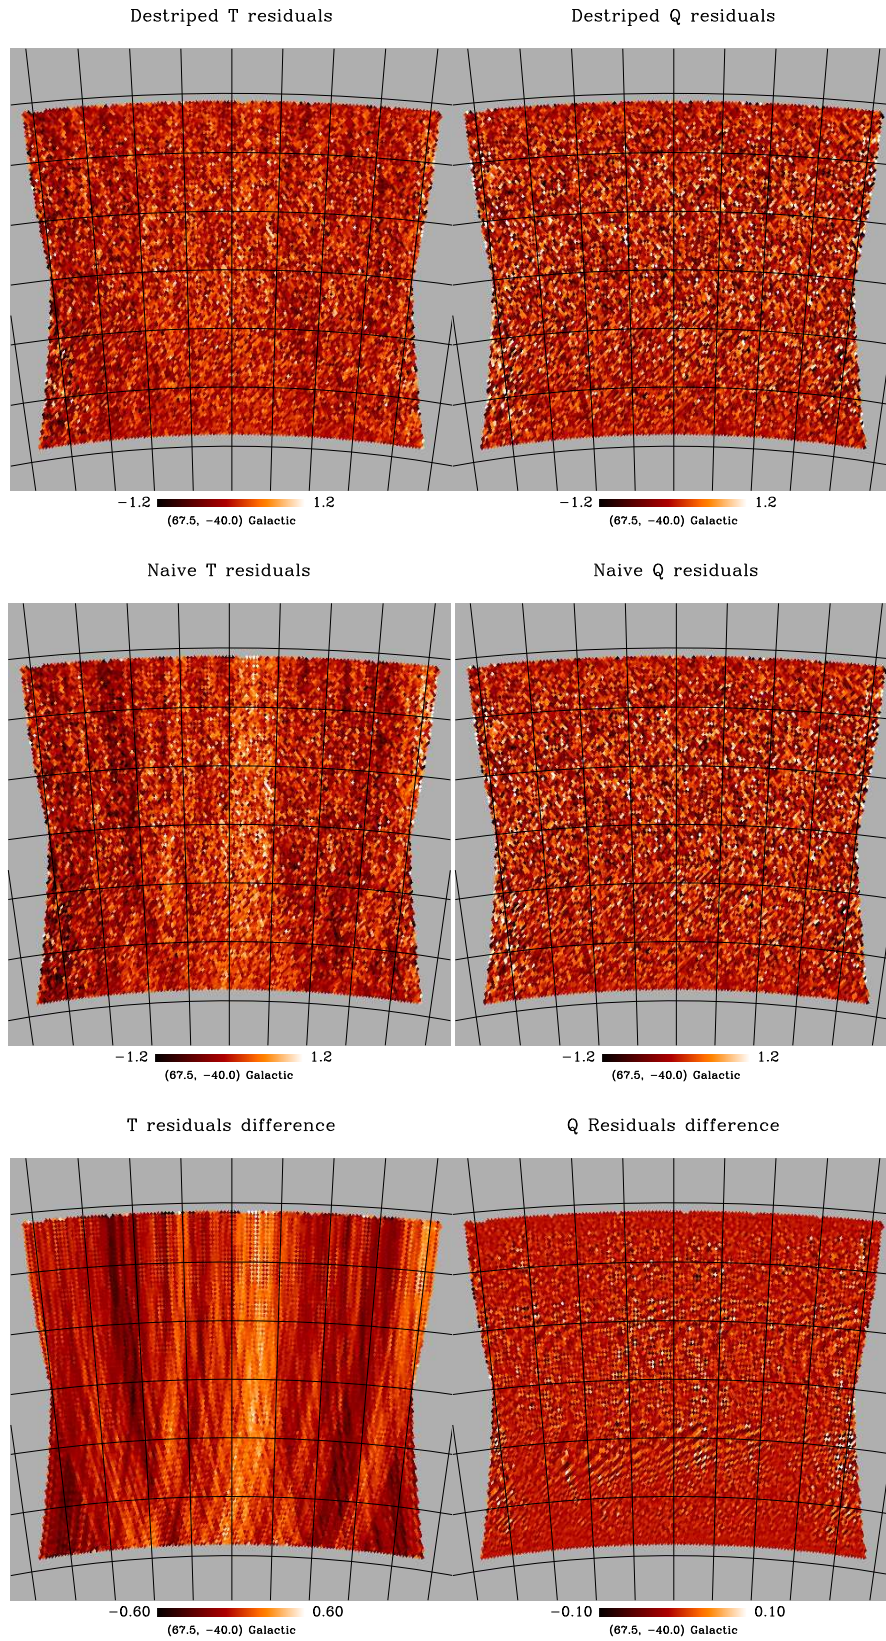


Figure 5. *Top row:* Residuals between a Descart 1-s output map and the input map, T on the left and Q on the right. *Middle row:* Residuals between a naive output map and the input map, T on the left and Q on the right. *Bottom row:* Difference between the top and middle row: Naive residuals - Descart residuals. Again, T is on the left and Q is on the right.




Figure 7. Variation of mean dimensionless Δ_Q with destriping baseline length λ_C in seconds, for the detector simulations.

maps have residuals nearly indistinguishable from optimality at all angular scales.

4.2.2 Signal and atmospheric noise simulations

The simulated atmospheric noise has more correlated noise, with $1/f_{knee} = 0.2$ -Hz and a much larger spectral index of $\alpha = 1.9$. Correspondingly, significantly more correlated noise leaks through to the T map from naive map-making. In this case, sophisticated mappers are essential, reducing the rms T residual by a factor of 10 and the Q residuals by a factor of 2 (Figure 9). The majority of the improvement is achievable using a baseline of 10-s, achieving 98.7% of the possible improvement in T residual power and 99.0% of the possible improvement in Q residual power. The best 1-s baseline maps achieve 99.9% of the possible improvement for T and 99.9% of the possible improvement for Q.

Despite the hardware effort of QU modulation using a half-wave plate, this noise regime produces reducible Q and U pixel noise sourced from correlated TOD noise that can be more than halved by the application of even relatively long baseline destriping. This assumes no $T \rightarrow P$ mixing from instrumental polarisation, the results of which would alias stripes into the Q and U maps themselves, strengthening the requirement for destriping.

The lower panel of Figure 8 shows the mean angular power spectra of the T residuals for the algorithms over 200 realisations (the curves are the same as those in the upper panel of Figure 8). It should be noted that the optimal map shows slightly more noise power for these simulations than it did in the detector simulations, because even the optimal maximum-likelihood solution cannot remove all of the $1/f$ noise.

The “quick and dirty” 400-s baseline approach is more effective at smaller angular scales than it was for the detector noise. The best 1-s baseline Descart maps are again nearly indistinguishable from optimal maps at all angular scales.

4.2.3 Fence scan simulations

The fence scan simulations used the same noise parameters as the signal and detector noise simulations with $f_{knee} = 0.1$ -Hz and

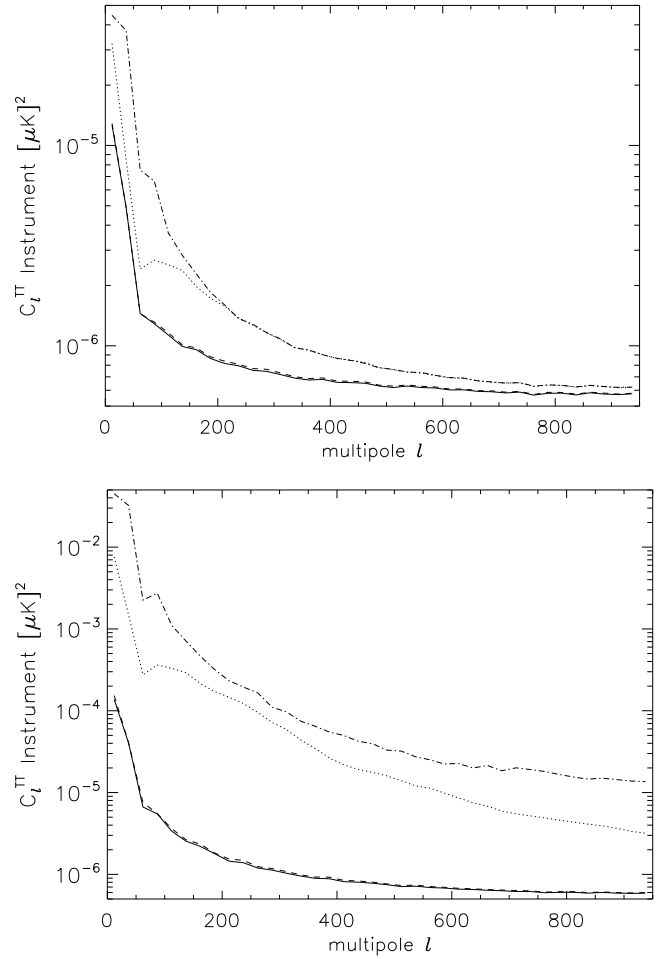


Figure 8. Upper panel: Mean T residual angular power spectra for the detector $f_{knee} = 0.1$ -Hz simulations. The solid curve is for the MLE residuals; the dashed curve is for the 1 second destriping residuals; the dotted curve is for the 400 seconds destriping residuals; the dot-dashed curve is for the naive binning (equation 2) residuals. Lower panel: Same as above, but for the mean T residual angular power spectra of the signal+atmospheric noise simulations.

$\alpha = 1.0$. The scan was designed to cover approximately the same number of pixels using the same number of observations. The average integration time per pixel is the same, however the isotropy and connectivity of the scans are different. Both the fence and sabre scans are sinusoidal scans, spending disproportionately more time integrating at the edges of the field, the fence scan integrating more at all four field edges and the sabre scan integrating more only on two edges. The connectivity, or level of cross-linking, is much greater in the fence scan, with the two scanning directions ideally cross-linked in being perpendicular to one another.

The rms residuals in the optimal fence maps are 8.7% smaller than the optimal sabre maps. Further to this, Figure 10 shows that the Descart T residuals decay toward optimality at longer baselines than for the sabre scan, reaching a trough at 10-s baselines: further reductions yield negligibly different residuals as the maps are already near-optimal.

The improvements in Q residuals are similar to the improvements in the sabre scan simulations: the level of cross-linking has




Figure 9. *Upper panel:* Variation of mean dimensionless Δ_T for the signal + atmospheric noise simulations. *Lower panel:* Variation of mean dimensionless Δ_Q for the signal + atmospheric noise simulations.

little effect on the near-white noise in the modulated Q and U streams.

4.3 E and B mode errors

The polarisation modulation design is intended to remove the effects of correlated noise from the Q and U maps and thence from the E and B power spectra. However, the increased variance in the Q and U maps from simple naive binning of a single detector propagates through to the polarisation power spectra.

Figures 11 and 12 show the mean estimated E and B-mode power spectra from the 200 signal+detector noise simulations respectively, after subtraction of the mean noise power spectrum. The error bars are the Monte-Carlo error bars for a single signal+noise realisation, and the dashed curves are the input theoretical power spectra. The dotted curve indicates the magnitude of the plotted error bar in each bin, which include both experimental noise and sample variance from partial sky coverage. The mean estimates are unbiased with respect to the input model.

The mean noise angular power spectra from the signal+detector noise simulations are shown in Figure 13. The curves are the difference between the noise angular power spectra ei-




Figure 10. *Upper panel:* Variation of dimensionless δ_T for the fence scan simulations. *Lower panel:* Variation of dimensionless δ_Q for the fence scan simulations.

ther from naive (solid curve) or destriping (dashed curve) and that of the maximum-likelihood noise power spectrum, defined by: $\langle N_l \rangle^{approx} - \langle N_l \rangle^{MLE}$, where $\langle N_l \rangle^{approx}$ is the noise spectrum from either the 1-s Descart maps or the naive maps. If either is optimal, the curve will be all zeros.

Whilst neither curve is optimal, the change in noise power is orders of magnitude smaller than the signal power at all multipoles and has a negligible effect on the error bar magnitudes.

Figure 14 shows the noise angular power spectra difference for the signal+atmospheric noise simulations, where the changes are very significant. At multipoles $l > 400$, the B-mode noise angular power spectrum becomes larger than the input signal angular power spectrum.

The effect of this on the B-mode error bars is clearly seen in the significance of the total detection. The total significance estimator \hat{C} bins the bandpowers of the spectrum (\hat{C}_b) into a single bin, weighted proportionally to the power of the input fiducial model C_b^{fid} and inversely proportional to the variance of the bandpower σ_b^2




Figure 11. Plot of the mean estimated C_l^{EE} for the $f_{knee} = 0.1$ -Hz simulations, using Descart on 200 noise and 200 signal + noise realisations with 1-s baselines. The error bars are the average for each individual signal+noise realisation. The dashed curve is the input power spectrum and the dotted curve is the magnitude of the plotted error bars.




Figure 12. As Figure 11 but for C_l^{BB} estimates from Descart 1-s maps.

$$\hat{C} = \frac{\sum_b \frac{C_b^{fid}}{\sigma_b^2} \hat{C}_b}{\sum_b \frac{(C_b^{fid})^2}{\sigma_b^2}}. \quad (43)$$

The significance of the detection is then given by $\frac{\langle \hat{C} \rangle}{\sqrt{\langle (\hat{C} - \langle \hat{C} \rangle)^2 \rangle}}$.

The significance of the E and B-mode detections for the algorithms with atmospheric noise are shown in Table 3. The increased noise from the naive maps slightly erodes the significance of the E-mode detection and destroys the detection of B-modes. We reiterate that this applies only in the case of a single detector and would not apply, for example, to detector differencing experiments.

4.4 Spurious B modes

200 signal+noise simulations were created using input CMB maps with artificially zero B-modes and typical detector $1/f$ ($f_{knee} =$





Figure 13. Upper panel: Differences between E-mode noise power spectra from the $f_{knee} = 0.1$ -Hz simulations in μK^2 . The solid curve is $\langle N_l^E \rangle_{naive} - \langle N_l^E \rangle_{MLE}$ whilst the dashed curve is $\langle N_l^E \rangle_{destriped} - \langle N_l^E \rangle_{MLE}$ for Descart with 1-s baselines. Lower panel: Differences between B-mode noise power spectra from the $f_{knee} = 0.1$ -Hz simulations in μK^2 . The solid curve is $\langle N_l^B \rangle_{naive} - \langle N_l^B \rangle_{MLE}$ whilst the dashed curve is $\langle N_l^B \rangle_{destriped} - \langle N_l^B \rangle_{MLE}$ for Descart with 1-s baselines.

Table 3. Atmospheric noise simulation ($f_{knee} = 0.2$, $\alpha = 1.9$), detection significance

	E-mode	B-mode
MLE	37.90	8.22
Descart 1-s	37.90	8.22
Naive	35.29	0.02

0.1-Hz). Maps were made from these TOD streams using each algorithm and the B-modes of the polarisation fields were estimated as in the previous section.

There was no evidence of the presence of spurious B-modes due to mode mixing from any of the algorithms. Figure 15 shows the mean estimated B mode for these TOD for Descart with 1-s baselines (the plots are identical for the naive mapper and the MLE algorithm). It is unbiased, correctly returning an ensemble average zero B mode estimate.




Figure 14. Upper panel: Same as the upper panel of Figure 13, but for the atmospheric noise simulations. Lower panel: Same as the lower panel of Figure 13, but for the atmospheric noise simulations.




Figure 15. As Figure 12 but for the simulation with zero C_l^{BB} from Descart 1-s maps.




Figure 16. Comparative runtime (t/t_{MLE}) variation with destriping length λ_c . Upper panel: runtimes for the signal+detector noise simulations. Lower panel: runtimes for the signal+atmospheric noise simulations.

4.5 Computing resources

The principal concern for computing resources for map-making is CPU time, as it is feasible to hold many single days' worth of TOD in memory at one time. The naive algorithm represents the quickest, and dirtiest, possible method for reducing TOD into a map. Its computation time is dominated by the input/output overhead.

The iterative scaling of the destriping algorithm is $\mathcal{O}(N_t + n_a \log_2 |n_a|)$, compared to $\mathcal{O}(N_t(1 + \log_2 |N_t|))$, for the maximum likelihood algorithm (where N_t is the number of number of observations). Effective values of n_a always satisfy $n_a \ll N_t$, for example, the n_a used in this paper are between 2 and 5 orders of magnitude smaller than N_t , so that the iterative scaling of Descart $N_t + n_a \log_2 |n_a| \approx N_t$. Noting that typical time stream lengths for day strategies are of order 10^7 , we can predict that destriping will require approximately an order of magnitude less CPU time to run.

In practice, the time-stream can be split into chunks to speed up the Fourier de-convolutions, where the chunk size is some multiple of the correlation length of the time-stream noise. For a chunk size of n_c the iterative scaling of the FFTs reduces to $\mathcal{O}(N_{tod} \log_2 |n_c|)$ for the maximum-likelihood algorithm and $\mathcal{O}(N_t + n_a \log_2 |n_{ca}|)$ for destriping algorithm, where $n_{ca} = n_c/\lambda_c$ is still 2 – 5 orders of magnitude smaller than n_c , maintaining the reduced complexity of the destriping algorithm. We also note that the use of hardware optimised FFT libraries, as opposed

Table 4. Typical numbers of Descart iterations for the signal+detector noise, signal+atmospheric noise and fence scan simulations. Also shown is the destriping baseline length λ_C in seconds and the corresponding number of offset functions in the system (n_a).

λ_C (s)	n_a	detector noise	atmospheric noise	fence scan
1	43200	40	127	18
5	8640	36	89	15
10	4320	32	65	15
100	432	22	25	11
400	108	16	19	9
1200	36	10	11	7

to the ubiquitous FFTW library, will further accelerate the FFT procedures.

Scaled computing times (t/t_{MLE} , where t_{MLE} is the CPU time required for the maximum likelihood algorithm to complete) for the algorithms are shown in Figure 16 for both the signal+detector noise simulations and the signal+atmospheric noise simulations. The codes are all serial and the runs were completed on the same machine, using dual-core Intel Xeon 3GHz CPUs, with the same allocated resources. The longest Descart baseline used was 1200-s, which returned a 49.1% reduction in Q and U residuals for the atmospheric noise simulations (Figure 9, lower panel), out of a possible 55.04% for the optimal algorithm, whilst achieving a $22\times$ speed improvement over the optimal algorithm.

In addition to the iterative scaling, the computing time is also determined by the number of iterations required to solve the system. Table 4 shows the typical numbers of iterations required by Descart for the simulations compared to the destriping baseline length used. The fence scan simulations required fewer iterations than the sabre scan simulations with the same noise parameters (detector noise column). The greater cross-linking of the fence scan has produced a system that is easier to solve than the sabre scan especially at short baseline lengths where the system (n_a) becomes large.

The number of iterations is similar for both sets of sabre scan simulations (detector noise and atmospheric noise columns) at very long baselines, but for the atmospheric noise increases much more quickly with baseline lengths < 100 -s, requiring $\approx 3\times$ as many iterations than the detector noise simulations at the near optimal 1-s baseline.

By modelling the correlated noise as a series of offset functions, the destriping algorithm is solving a smaller system than the full maximum likelihood algorithm (generally $n_a < n_p$ except for the very shortest baselines, where n_a and n_p are number of offsets and pixels respectively). Typical experimental resolutions in future experiments will be higher than that used here ($n_{side} = 512$), for example $n_{side} = 2048$ for C_l OVER, increasing n_p by a factor of 16 and leaving n_a untouched.

4.6 Importance of including C_a

Prior to the development of the MADAM algorithm (Keihänen et al. 2005), destriping was conducted without prior noise information. This “traditional” destriping used only the *data term* in (28) (Burigana et al. (1997), Delabrouille (1998), Maino et al. (1999) Keihänen et al. (2004)), neglecting the *prior term* containing C_a . Without this noise information, the potential of destriping to remove correlated noise is considerably reduced. It




Figure 17. Comparison between destriping with (solid line with circles) and without (dashed line with \times) C_a information.

is not possible to use traditional destriping at λ_C s where there is significant correlations between the offset function amplitudes a_i .

Figure 17 shows the effects of using traditional destriping in this regime. The dashed line shows the returned destriping Δ_T vs λ_C . At large λ_C , where the offsets are uncorrelated, traditional destriping performs similarly to the covariant destriping using C_a (the solid curve with circles). When λ_C is reduced into the regime of correlated offsets, the residuals behave pathologically, exploding to considerably higher magnitude than for the naive maps.

This can be explained in two ways. The noise model in the data term of (28) assumes no correlations between the offsets. When correlations are present, the noise model breaks down and the results are junk. Including C_a introduces prior information on noise correlations into the model. This can also be recognised as linear regularisation, where the introduction of the regularising matrix C_a forces the solution to the system to be well behaved.

This highlights the vital importance of including C_a . With its inclusion, destriping can be used to its full potential in returning near optimal maps at a fraction of the computing time.

5 DISCUSSION

In this paper, we set out to compare the destriping approach to map-making to the optimal maximum likelihood estimator. In particular, we want to know if destriping is fast and accurate enough to replace the maximum likelihood estimator for partial-sky, polarisation modulation experiments.

We have written a new covariant destriping code named Descart (DEStriping CARTographer) and a new maximum-likelihood estimation code and applied them to non-circular scanning strategies for the first time. We introduce variations in cross-linking through our realistic, but non-ideal, sabre scan and the unrealistic, but ideally cross-linked, fence scan.

Neither of these strategies allow the co-addition of circular scans before destriping, a process relied upon by the early “traditional” destriping efforts (Burigana et al. 1997, Delabrouille 1998, Maino et al. 1999, Keihänen et al. 2004). We must use a “covariant” destriping algorithm (Keihänen et al. 2005), incorporating prior knowledge of the noise power spectrum through C_a . This al-

lows the use of very short baselines that produce nearly optimal residuals in the output T, Q and U maps. We have shown that, without this prior, the noise model clearly breaks down pathologically at shorter baselines and the returned map becomes junk.

The simulations we produced applied destriping to different noise regimes than those to which it has previously been applied. We have probed regimes of both detector noise and atmospheric noise, the latter of which represents a much greater challenge as it dominates noise even in the current generation of ground-based polarisation experiments. Previous work on the efficacy of destriping concentrated on the Planck satellite mission, assuming low levels of correlation in the $1/f$ noise, as is possible in the absence of the atmosphere, using knee frequencies of 0.03-Hz for the HFI (Ashdown et al. 2007a) and 0.05-Hz for the LFI (Ashdown et al. 2007b). We have simulated detector noise with $f_{knee} = 0.1$ -Hz and atmospheric noise with $f_{knee} = 0.2$ -Hz and a spectral index $\alpha = 1.9$, similar to the level experienced by QuAD (Hinderks et al. 2008).

Destriping performs better in the more realistic and challenging atmospheric noise simulations than in the detector only noise simulations. For the shortest baseline of 1-s, it achieves 99.9% of the possible improvement in T, Q and U residuals achievable with the full maximum-likelihood algorithm. This amounts to reduction in T residual magnitude of 90.02% and in Q and U of 55.04%. One of the motivators of this study is the requirement to produce high quality T maps, in addition to the Q and U maps, in order to remove $T \rightarrow P$ leakage from instrumental polarisation (eg: Johnson et al. 2007), which will also leak atmospheric $1/f$ into the Q and U signal despite the modulation of those signal by the rotating half-wave plate.

We have found that even without the simulation of this leakage, it is vital to deal with the correlated noise in order to observe B-modes: with naive mapping, the significance of the B-mode detection was destroyed, falling from 8.22 σ for the optimal and 1-s Descart algorithms to 0.02 σ . This is a conclusion specific to the single detector case, as the atmospheric noise is completely correlated between detector pairs that share the same horn. Exploiting this correlation, by combining data from both channels, can completely remove the atmospheric noise.

The detector noise simulations maps also saw significant improvement, with the 1-s baseline Descart maps showing 96% of the possible improvement in T and 94.71% of the possible improvement in Q residuals, though the latter amounts to less than 0.1% of the noise power in this case. Whilst this is not reducible by detector time-stream differencing, the effect on the B-mode noise was very small.

The 1-s Descart maps were returned with improvements in computing time of a factor of 6 for both detector and atmospheric noise. However, for the atmospheric noise simulations, the 98.73% of the improvement in T and 99.9% of Q and U was returned using a longer baseline of 10-s, which delivered an improvement in speed of 10 \times . Most significantly, 89.16% of the improvement in Q residuals were achieved with a very long baseline of 1200-s, returning a speed improvement of 22 \times over the maximum likelihood algorithm. This can be understood by noting that the atmospheric noise has a lot more noise power at low frequencies than the detector noise simulations do. As the resolution of C_a is increased, improvements in the noise are reached more rapidly as the extra noise is resolved.

The destriping algorithm was shown not to distort the CMB signal, which is the prohibitive drawback of fast filtering methods, such as the $C_N^{-1/2}$ filter, which introduces extra variance into the

power spectra. Further, destriping is unbiased for B-modes, producing correct null detections for simulations with artificially zero B-modes.

The ideally cross-linked fence scan showed considerably quicker reduction in residuals for T with decreasing baseline length, than did the sabre scan with detector noise, despite identical noise in the two sets of simulations. Improvement in residuals reached a trough at 10-s baselines: further improvements were negligible for shorter baselines. Q and U residuals improved similarly to the sabre scan.

Covariant destriping produces very close to optimal maps, but much more quickly than the full maximum-likelihood algorithm, at speeds that make it applicable to the large datasets from upcoming B-mode experiments. Destriping can also reduce the vast majority of atmospheric noise in single detector time-streams very quickly, using more approximate long baselines.

This paper is the first in a planned series investigating destriping for ground based polarisation experiments. The focus of future work will be for experiments in which the Q and U signals are subject to $1/f$ noise (eg: QUIET). For these experiments, the reduction in $1/f$ will be critically important to the science return. The simulations will be extended to higher resolution and will investigate the effects of sub-pixel signal gradients by using CMB input maps at much higher resolution than the recovered maps. The simulations will also be extended to multiple detectors, with noise correlated between detectors sourced both from detector fluctuations and from the atmosphere. Destriping will be applied alongside and directly compared to the only realistic alternative: the Monte-Carlo based MASTER method where the time-stream is high-pass filtered prior to map-making (Hivon et al. 2002).

Acknowledgments: We are grateful to Hans Kristian Eriksen, Anthony Challinor, Torsti Poutanen, Chris North, Joe Zuntz, Giancarlo de Gasperis and the C_ℓ OVER collaboration for useful discussions, and to Jonathan Patterson for invaluable computing support. We acknowledge the use of the HEALPix and FFTW subroutine libraries. David Sutton acknowledges the support of a Science and Technology Facilities Council PhD Studentship.

REFERENCES

- Amblard A., Hamilton J-Ch., 2004, A&A, 417, 1189
- Ashdown M. A. J. et al., 2007a, A&A, 467, 761
- Ashdown M. A. J. et al., 2007b, A&A, 471, 361
- Barret R. et al., 2006, Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods, 2nd Edition, SIAM
- Borrill J., 1999, pre-print (astro-ph/9911389v1)
- Brown M. L., Castro P. G., Taylor A. N., 2005, MNRAS, 360, 1262
- Bunn E. F., 2002, Phys. Rev. D, 65, 043003
- Burigana C., Malaspina M., Mandolesi N., Danese L., Maino D., Bersanelli M., Maltoni M., 1997, pre-print (astro-ph/9906360)
- Bussmann R. S., Holzapfel W. L., Kuo C. L., ApJ, 622, 1343
- Challinor A., Chon G., 2005, MNRAS, 360, 509
- de Gasperis G., Balbi A., Cabella P., Natoli P., Vittorio N., 2005, A&A, 436, 1159
- Delabrouille J., 1998, A&AS, 127, 555
- Dore O., Teyssier R., Bouchet F.R., Vibert D., Prunet S., 2001, A&A, 374, 358
- Ferreira P. G., & Jaffe A. H., 2000, MNRAS, 312, 89
- Gorski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759

- Hanany S., Rosenkranz P., 2003, NewAR, 47, 1159
 Hanany S. et al. 2000, ApJ, 545, L5
 Hinderks J. et al., 2008, pre-print (astro-ph/0805.1990v1)
 Hinshaw G. et al. 2008, pre-print (astro-ph/0803.0732)
 Hivon E., Gorski K. M., Netterfield C. B., Crill P. C., Prunet S., Hansen F., 2002, ApJ, 567, 2
 Johnson B. R. et al., 2007, ApJ, 665, 42J
 Keihanen E., Kurki-Suonio H., Poutanen T., Maino D., Burigana C., 2004 A&A, 428, 287
 Keihanen E., Kurki-Suonio H., Poutanen T., 2005, MNRAS, 360, 390
 Lewis A., Challinor A., Lasenby A., 2000, ApJ, 538, 473
 Maino D. et al., 1999, A&AS, 140, 383
 MacTavish C. J. et al., 2007, pre-print (astro-ph/0710.0375)
 Macias-Perez J. F. et al., 1999, A&A, 467, 1313
 Masi S. et al., 2006, ApJ, 458, 687
 Natoli P., de Gasperis G., Gheller C., Vittorio N., 2007, A&A, 372, 346
 North C. E. et al., 2008, pre-print (astro-ph/0805.3690)
 Oxley P. et al., 2005, pre-print (astro-ph/0501111)
 Patanchon G. et al., 2007, pre-print (astro-ph/0711.3462)
 Poutanen et al., 2006, A&A, 449, 1311
 Press et al., 2002, Numerical Recipes in C++, 2nd Edition, Cambridge
 Smith K. M., 2006, Phys. Rev. D, 74, 083002
 Smith K. M., Zaldarriaga M., 2007, Phys. Rev. D, 76, 043001
 Smoot G. F. et al., 1992, ApJ, 396, L1
 Stompor R. et al., 2002, Phys. Rev. D, 65, 022003
 Tegmark M., 1997a, ApJ, 480, L87-L90
 Tegmark M., 1997b, Phys. Rev. D, 56, 84514
 Zaldarriaga M., Seljak U., 1997, Phys. Rev. D, 55, 4 1830

APPENDIX A: PURE PSEUDO- C_L ESTIMATION

Spatial information about the comparative map-making residuals is gained by analysing the angular power spectrum of the residual maps. Estimating B-mode power spectra with partial sky coverage is a nontrivial problem, due largely to the presence of ambiguous modes (Bunn 2002; Challinor & Chon 2005) which receive contributions from both E-mode and B-mode power on the full sky. This “ $E \rightarrow B$ mixing” can act as an extra source of noise, for a B-mode power spectrum estimator which does not filter out ambiguous modes. Because B-modes are a primary scientific target for the next-generation polarisation experiments considered in this paper, it will be important for the analysis pipelines of such experiments to use a B-mode power spectrum estimator which does not suffer from $E \rightarrow B$ mixing. Accordingly, we analyse our residual maps using one such estimator: the pure pseudo- C_ℓ estimator from (Smith 2006; Smith & Zaldarriaga 2007), which generalizes the MASTER construction (Hivon et al. 2002; Brown, Castro & Taylor 2005) by eliminating $E \rightarrow B$ mixing from ambiguous modes. In this appendix, we briefly summarize its construction and key properties.

First, in order to apply pure pseudo- C_ℓ power spectrum estimation, one must choose (heuristically) a pixel weight function $W(\hat{n})$. In this paper we have used cosine apodization: in spherical coordinates with north pole at the center of the survey, the weight function is given by:

$$W(\theta, \phi) = \begin{cases} 1 & \theta < r - r_* \\ \frac{1}{2} - \frac{1}{2} \cos(\pi \frac{r-\theta}{r_*}) & r - r_* \leq \theta \leq r \\ 0 & \theta > r \end{cases} \quad (\text{A1})$$

where r is the survey radius and r_* is an apodization length. In this paper, the survey radius is $r = 7^\circ$ and heuristic apodisation radii of $r_* = 4.3^\circ$ and $r_* = 1.61^\circ$ are used for multipole ranges $\ell \leq 40$ and $\ell > 40$ respectively.

We then define pseudo multipoles $\tilde{a}_{\ell m}^E, \tilde{a}_{\ell m}^B$ by:

$$\tilde{a}_{\ell m}^E = -\frac{1}{2} \sqrt{\frac{(l-2)!}{(l+2)!}} \int d^2 \hat{n} \left[\Pi_+(\hat{n}) W(\hat{n}) \bar{\partial} \bar{\partial} Y_{\ell m}^*(\hat{n}) + \Pi_-(\hat{n}) W(\hat{n}) \bar{\partial} \bar{\partial} Y_{\ell m}^*(x) \right] \quad (\text{A2})$$

$$\tilde{a}_{\ell m}^B = -\frac{i}{2} \sqrt{\frac{(l-2)!}{(l+2)!}} \int d^2 \hat{n} \left[\Pi_+(\hat{n}) \bar{\partial} \bar{\partial} (W(\hat{n}) Y_{\ell m}^*(\hat{n})) - \Pi_-(\hat{n}) \bar{\partial} \bar{\partial} (W(\hat{n}) Y_{\ell m}^*(x)) \right] \quad (\text{A3})$$

where $\Pi_\pm(\hat{n}) = (Q \pm iU)(\hat{n})$ and $\bar{\partial}, \bar{\partial}$ are the spin raising and lowering operators defined in Zaldarriaga, & Seljak (1997). The derivative operators have been placed differently in the definitions of $\tilde{a}_{\ell m}^E, \tilde{a}_{\ell m}^B$ above, so that $\tilde{a}_{\ell m}^E$ will receive contributions from ambiguous modes, but $\tilde{a}_{\ell m}^B$ will not. Note that Π_\pm is a spin (± 2) field, but the integrands appearing in Eqs. (A2), (A3) are spin-zero and do not depend on the choice of local frame.

The next step is to define pseudo power spectra, in bands b , by:

$$\tilde{C}_b^{EE} = \frac{1}{2\ell+1} \sum_{\ell m} P_{b\ell} |\tilde{a}_{\ell m}^E|^2 \quad (\text{A4})$$

$$\tilde{C}_b^{BB} = \frac{1}{2\ell+1} \sum_{\ell m} P_{b\ell} |\tilde{a}_{\ell m}^B|^2 \quad (\text{A5})$$

where $P_{b\ell}$ is a binning operator ($P_{b\ell} = 0$ unless $\ell \in b$).

As constructed, the pseudo spectra are biased estimators of the signal power spectra; one has

$$\begin{pmatrix} \langle \tilde{C}_b^{EE} \rangle \\ \langle \tilde{C}_b^{BB} \rangle \end{pmatrix} = \begin{pmatrix} K_{bb'}^+ & K_{bb'}^- \\ 0 & K_{bb'}^{+pure} \end{pmatrix} \begin{pmatrix} C_{b'}^{EE} \\ C_{b'}^{BB} \end{pmatrix} \quad (\text{A6})$$

where the transfer matrices $K_{bb'}$ depend only on the survey geometry and pixel weighting, and can be computed efficiently using the algorithm in (Smith 2006). Note that the lower left block in Eq. (A6) is zero because $\tilde{a}_{\ell m}^B$ does not receive contributions from ambiguous modes. We do not include an additive term from noise bias on the right-hand side of Eq. (A6) as this equation is for the case of estimating the power spectrum of a pure noise map.

The final step in the construction is to define unbiased estimators $\hat{C}_b^{EE}, \hat{C}_b^{BB}$ by:

$$\begin{pmatrix} \hat{C}_b^{EE} \\ \hat{C}_b^{BB} \end{pmatrix} = \begin{pmatrix} K_{bb'}^+ & K_{bb'}^- \\ 0 & K_{bb'}^{+pure} \end{pmatrix}^{-1} \begin{pmatrix} \tilde{C}_b^{EE} \\ \tilde{C}_b^{BB} \end{pmatrix} \quad (\text{A7})$$

These are unbiased ($\langle \hat{C}_b^{EE} \rangle = C_b^{EE}$ and $\langle \hat{C}_b^{BB} \rangle = C_b^{BB}$) power spectrum estimators, with no $E \rightarrow B$ mixing: the B-mode estimator has the property that it receives no contributions from ambiguous modes.

APPENDIX B: DESTRIPIING FOR MULTIPLE DETECTORS INCLUDING CORRELATIONS

The time ordered output from each detector, y_t , is stacked end to end to form a single $N_{tod} \times N_{detector}$ vector. The vector of

offset function amplitudes a_α , where α indexes the offset function, is similarly stacked end to end to form a $N_{offsets} \times N_{detector}$ vector, as are the projection matrices P_{tp} and $F_{t\alpha}$.

Each time ordered datum is now indexed by $i = lt$, where l denotes a particular detector and t denotes a particular time. Likewise, we index the offsets as $j = l\alpha$. With the new indices, the full focal plane TOD vector y_i is modelled as

$$y_i = P_{ip}x_p + F_{ij}a_j + n_{W,i} \quad (B1)$$

where $F_{ij}a_j$ models the correlated noise component of y_i and $n_{W,i}$ is the uncorrelated white noise.

The likelihood for the system to be minimised is

$$-2 \ln L = (y_i - P_{ip}x_p - F_{ij}a_j)^T C_{W,ii'}^{-1} (y_{i'} - P_{i'p}x_p - F_{i'j'}a_{j'}) + a_j^T C_{a,jj'}^{-1} a_{j'} \quad (B2)$$

The white noise covariance matrix $C_{W,ii'}$ is n_{TOD} independent $N_{detector}^2$ matrices describing the correlation of the white noise between detectors. For atmospheric common mode noise, we assume that this matrix is diagonal

$$C_{W,ii'} = \sigma^2 \delta_{ii'} \quad (B3)$$

so that white noise is uncorrelated between detectors.

The offset covariance matrix C_a encodes prior information on the correlations in the long term atmospheric $1/f$ between detectors. Each $(N_{offset} \times N_{offset})$ sub-matrix is not diagonal but circulant and can be inverted in the Fourier domain, using the method of Patanchon et al. (2007).

We define a multi-detector Fourier transform operator F , such that the Fourier transforms \tilde{y}_k of the TOD from each detector, stacked end to end, is

$$\tilde{y}_k = F y_i \quad (B4)$$

where the index $k = lf$ denotes the Fourier mode f of detector l . F is a block diagonal matrix, in which each block is Fourier transform operator for a single detector channel.

The Fourier domain offset covariance matrix is

$$R_{kk'} = F C_{a,jj'} F^\dagger. \quad (B5)$$

Each detector-detector sub-matrix of R , which we label $R_{ll'}$ corresponding to detector combination ll' , is diagonal and each diagonal element is an independent Fourier mode. Its inverse R^{-1} is easy and quick to compute explicitly, an operation that only needs to be accomplished once.

The inverse of the offset covariance is obtained by

$$[C_a^{-1}]_{ll'} = F^\dagger [R^{-1}]_{ll'} F. \quad (B6)$$

However, there is no need to calculate and store this in the time domain, as the operation $C_a^{-1}a$ is completed more quickly by switching between time and Fourier space

$$C_a^{-1}a = \mathcal{F}^{-1}[R^{-1}\mathcal{F}[a]]. \quad (B7)$$

Armed with this technique for painlessly inverting the offset covariance matrix, we can build the estimator for the offset amplitudes

$$(\mathbf{F}_{j'i'} \mathbf{Z}_{i'i} \mathbf{F}_{ij} + \sigma_w^2 \mathbf{C}_{a,jj'}^{-1}) a_j = \mathbf{F}_{j'i'} \mathbf{Z}_{i'i} y_i \quad (B8)$$

where $\mathbf{Z}_{ii'}$ is block-diagonal in l and each detector's block diagonal sub-matrix is given by (30).